Quantitative Information Flow for Scheduler-Dependent Systems *

Yusuke Kawamoto
Inria Saclay & LIX, École Polytechnique, France

Thomas Given-Wilson
Inria, France

Quantitative information flow analyses measure how much information on secrets is leaked by publicly observable outputs. One area of interest is to quantify and estimate the information leakage of composed systems. Prior work has focused on running disjoint component systems in parallel and reasoning about the leakage compositionally, but has not explored how the component systems are run in parallel or how the leakage of composed systems can be minimised. In this paper we consider the manner in which parallel systems can be combined or scheduled. This considers the effects of scheduling channels where resources may be shared, or whether the outputs may be incrementally observed. We also generalise the attacker's capability, of observing outputs of the system, to consider attackers who may be imperfect in their observations, e.g. when outputs may be confused with one another, or when assessing the time taken for an output to appear. Our main contribution is to present how scheduling and observation effect information leakage properties. In particular, that scheduling can hide some leaked information from perfect observers, while some scheduling may reveal secret information that is hidden to imperfect observers. In addition we present an algorithm to construct a scheduler that minimises the min-entropy leakage and min-capacity in the presence of any observer.

1 Introduction

Preventing the leakage of confidential information is an important goal in research of information security. When some information leakage is unavoidable in practice, the next step is to quantify and reduce the leakage. Recently theories and tools on quantitative information flow have been developed using information theory to address these issues [17, 7, 23, 12, 25, 8, 16, 15]. The common approach is to model systems as *information-theoretic channels* that receive secret input and returns observable output.

One area of interest is to quantify and estimate the information leakage of composed systems. When composing systems the manner of reasoning about their behaviour is non-trivial and is complicated by many factors. One of the first approaches is to consider the *(disjoint) parallel composition*, that is, simply running the component systems independently and regarding them as a single composed system. This approach provides some general behaviour and reasoning about the whole composed system, as shown in the research of quantitative information flow with different operational scenarios of attack [5, 19, 20]. However, the parallel composition approach is coarse-grained and abstracts many of the channels' behaviours that may lead to changes in information leakage. Although this approach provides useful results on the bounds of possible leakage, it does so under the assumption that the component channels are executed independently and observed separately. That is, their outputs can always be linked to the disjoint component channels, and that both their outputs are observed simultaneously and without any interleaving or reflection of how the component channels achieved their outputs.

*This work has been partially supported by the project ANR-12-IS02-001 PACE, by the INRIA Equipe Associée PRINCESS, by the INRIA Large Scale Initiative CAPPRIS, and by EU grant agreement no. 295261 (MEALS). The work of Yusuke Kawamoto has been supported by a postdoc grant funded by the IDEX Digital Society project.

N. Bertrand and M. Tribastone (Eds.): QAPL 2015
EPTCS 194, 2015, pp. 48–62, doi:10.4204/EPTCS.194.4

© Y. Kawamoto & T. Given-Wilson
This work is licensed under the
Creative Commons Attribution License.

Here we take a more fine-grained approach where we consider that channels may provide a sequence of observable actions. Thus, a channel may be observed to output a sequence of actions, or the passage of time may be observed to pass between the initiation of the channel and a final output. This captures more mechanics of real world systems and allows for greater refined reasoning about their behaviour.

Such sequences of observable actions also allow a more subtle approach to combining channels in parallel. Rather than simply taking both outputs to appear together at the termination of their operations, observations can be made of the sequence in which the outputs appear. Such a combination of channels becomes parametrised by a *scheduler*, that informs on how to combine the observable sequences of actions into a single sequence. This can then represent very direct behaviour such as scheduling properties of a shared CPU, or abstract behaviours such as routing properties, vote counting, etc.

The other novel approach presented here is the refinement of the attacker's capability of observing the outputs of systems. We model attackers that may have imperfect observability: they may not accurately detect differences in outputs, or may do so only probabilistically. This captures, for example, the situation where the attacker may be blind to some internal behaviour that other agents can detect. In this paper such imperfect observations are modeled using what we call *observer channels*. This formalisation enables us to consider a large class of observers, including *probabilistic observers*, which have never been considered in the previous studies on quantitative information flow.

These refinements to composing information-theoretic channels allow us to reason about behaviours that may be obvious, but not captured by previous approaches. In this paper we present three kinds of results regarding the effect of leakage properties due to the considering of schedulers and observers. First, since scheduled composition can alter the leakage relative to the parallel composition, we present theorems for detecting when a scheduled composition does not alter the relative information leakage. This means some preliminary analysis may be sufficient to determine when scheduled composition may be worthy of further consideration. Second, scheduled composition can leak more or less information than the parallel composition depending on the properties of the channels and the power of the observer. Although the potential effect on leakage is dependent upon many factors, we present results that determine an upper bound for the leakage of a schedule-composed channel. Third, we present results for finding a scheduler that minimises the min-entropy leakage and min-capacity in the presence of any observer. We present how to construct such a scheduler by solving a linear programming problem.

In addition we evaluate our model and results with some simple yet intuitive examples, such as mix networks for voter anonymity, and side-channel attacks against cryptographic algorithms. We provide an implementation that can be used to calculate the behaviours of information-theoretic channels, schedulers, and observers as presented here. The implementation is available online [1], which requires the libraries leakiEst tool [14] and the linear programming system lp_solve [2].

The rest of the paper is structured as follows. Section 2 recalls the definitions of information-theoretic channels and measures of information leakage. Section 3 defines traces, systems and channel compositions, and shows examples of schedulers. Section 4 introduces the notion of generalised observers and defines the observed leakage. Section 5 presents our main results in a general manner. Section 6 applies these results to well known problems. Section 7 discusses some related work. Section 8 draws conclusions and discusses future work. All proofs can be found in [21].

2 Preliminaries

2.1 Information-Theoretic Channel

Systems are modeled as *information-theoretic channels* to quantify information leakage using information theory. A channel $\mathscr{K}$ is defined as a triple $(\mathscr{X},\mathscr{Y},C)$ consisting of a finite set $\mathscr{X}$ of secret input

values, a finite set $\mathscr{Y}$ of observable output values, and a *channel matrix* C each of whose row represents a probability distribution; i.e., for all $x \in \mathscr{X}$ and $y \in \mathscr{Y}$, $0 \leq C[x,y] \leq 1$ and $\sum_{y' \in \mathscr{Y}} C[x,y'] = 1$. For each $x \in \mathscr{X}$ and $y \in \mathscr{Y}$, $C[x,y]$ is a conditional probability $p(y|x)$ of observing y when the secret of the system is x. We assume some secret distribution π on $\mathscr{X}$, which is also called a *prior*. Given a prior π on $\mathscr{X}$, the joint distribution of having a secret $x \in \mathscr{X}$ and an observable $y \in \mathscr{Y}$ is defined by $p(x,y) = \pi[x]C[x,y]$.

2.2 Quantitative Information Leakage Measures

In this section we recall the definitions of two popular quantitative information leakage measures.

Mutual information is a leakage measure based on the Shannon entropy of the secret distribution.

Definition 1 Given a prior π on $\mathscr{X}$ and a channel $\mathscr{K} = (\mathscr{X}, \mathscr{Y}, C)$, the *mutual information* $\mathscr{I}(\pi, \mathscr{K})$ w.r.t. π and $\mathscr{K}$ is defined by:

$$\mathscr{I}(\pi, \mathscr{K}) = \sum_{x \in \mathscr{X}, y \in \mathscr{Y}} \pi[x]C[x,y] \log \left(\frac{C[x,y]}{\sum_{y' \in \mathscr{Y}} C[x,y']} \right).$$

Then the *Shannon's channel-capacity* $\mathscr{SC}(\mathscr{K})$ of a channel $\mathscr{K}$ is given by $\max_{\pi} \mathscr{I}(\pi, \mathscr{K})$ where π ranges over all distributions on $\mathscr{X}$.

Min-entropy leakage quantifies information leakage under single-attempt guessing attacks [9, 25].

Definition 2 Given a prior π on $\mathscr{X}$, and a channel $\mathscr{K} = (\mathscr{X}, \mathscr{Y}, C)$, the *prior vulnerability* $V(\pi)$ and the *posterior vulnerability* $V(\pi, \mathscr{K})$ are defined respectively as

$$V(\pi) = \max_{x \in \mathscr{X}} \pi[x] \quad \text{and} \quad V(\pi, \mathscr{K}) = \sum_{y \in \mathscr{Y}} \max_{x \in \mathscr{X}} \pi[x]C[x,y].$$

Then the *min-entropy leakage* $\mathscr{L}(\pi, \mathscr{K})$ and the *min-capacity* $\mathscr{MC}(\mathscr{K})$ are defined by:

$$\mathscr{L}(\pi, \mathscr{K}) = -\log V(\pi) + \log V(\pi, \mathscr{K}) \quad \text{and} \quad \mathscr{MC}(\mathscr{K}) = \sup_{\pi} \mathscr{L}(\pi, \mathscr{K}).$$

3 Information Leakage of Scheduler-Dependent Systems

3.1 Traces and Systems

In general the output of an information-theoretic channel can be defined in many different ways. In this work we consider traces, or sequences of actions, as observable values. Assume a countable set of *names* denoted $m, m', m_1, m_2, \ldots$ and a countable set of *values* $v, v_1, v', \ldots$. We define an *action* by $\mu, \alpha, \beta ::= \tau \mid \overline{m}\langle v \rangle$. Here τ denotes the traditional *silent* or *internal* action that contains no further information. The *output* action $\overline{m}\langle v \rangle$ can be considered to exhibit some value v via some named mechanism m. In concurrency theory the output action typically refers to the the named mechanism as a *channel name*, which is distinct from the notion of information-theoretic channel used here. Here the output action is used in a more general sense, in that $\overline{m}\langle v \rangle$ exhibits some value v such as runtime measured via mechanism m. For example, v could be runtime, electronic power usage or other value determined by the input, and m could be via direct communication/circuitry, indirect side effects, or any other means.

A *trace* is defined to be a sequence of actions of the form $\mu_1.\mu_2.\ldots.\mu_i$. The notation $\alpha \in \mu_1.\mu_2.\ldots.\mu_i$ denotes that there exists a $j \in \{1, 2, \ldots, i\}$ such that $\mu_j = \alpha$. Similarly a sequence of i actions μ can be denoted μ^i, and an empty sequence of actions by $\emptyset$. A *system* is modeled as an information-theoretic channel $(\mathscr{X}, \mathscr{Y}, C)$ where $|\mathscr{X}|$ is finite and the set $\mathscr{Y}$ of observables is a finite set of traces.

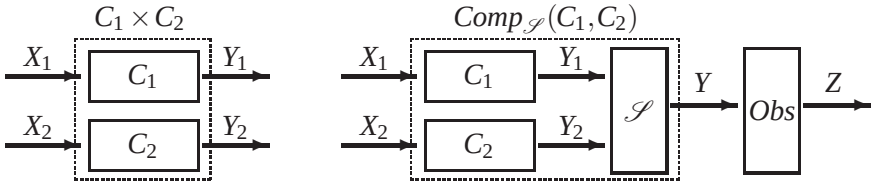

(a) Parallel composition (b) Observation of a scheduled composition

Figure 1: Parallel composition and scheduled composition

3.2 Scheduled Composition

In this section we model scheduler-dependent systems by introducing the notion of a *scheduled composition* of information-theoretic channels, which interleaves outputs from different channels.

In [20] the *parallel composition* $\mathscr{K}_1 \times \mathscr{K}_2$ of two component channels $\mathscr{K}_1$ and $\mathscr{K}_2$ is defined as a channel that outputs ordered pairs consisting of the outputs of the two component channels. That is, given two component channels $\mathscr{K}_1 = (\mathscr{X}_1, \mathscr{Y}_1, C_1)$ and $\mathscr{K}_2 = (\mathscr{X}_2, \mathscr{Y}_2, C_2)$, the outputs of their parallel composition range over the ordered pairs (y_1, y_2) for all $y_1 \in \mathscr{Y}_1$ and $y_2 \in \mathscr{Y}_2$. This composition can be modeled using a scheduler that allows $\mathscr{K}_1$ to perform the whole sequence y_1 of actions and some action $sep \notin \mathscr{Y}_1 \cup \mathscr{Y}_2$ (for separating y_1 from y_2) before $\mathscr{K}_2$ performs the actions in y_2.[1] In this setting we can recognise which component channel each output of the composed channel came out of.

In this paper we consider more fine-grained schedulers that may allow $\mathscr{K}_2$ to perform some actions before $\mathscr{K}_1$ completes the whole sequence of actions. To model such schedulers, we define the set of possible interleaving of two traces that preserves the orders of occurrences of actions in the traces.

Definition 3 (Interleaving of traces) Let us consider two traces y_1 of the form $\alpha_1.\alpha_2.\ldots.\alpha_k$ and y_2 of the form $\beta_1.\beta_2.\ldots.\beta_l$. The *interleaving* $Int(y_1, y_2)$ of y_1 and y_2 is the set of all traces of the form $\mu_1.\mu_2.\ldots.\mu_{k+l}$ s.t., for two sequences of distinct integers $1 \leq i_1 < i_2 < \ldots < i_k \leq k+l$ and $1 \leq j_1 < j_2 < \ldots < j_l \leq k+l$, we have $\mu_{i_m} = \alpha_m$ for all $m = 1, 2, \ldots, k$ and $\mu_{j_m} = \beta_m$ for all $m = 1, 2, \ldots, l$.

Definition 4 For two sets $\mathscr{Y}_1, \mathscr{Y}_2$ of observables, the *interleaving* $Int(\mathscr{Y}_1, \mathscr{Y}_2)$ *over* $\mathscr{Y}_1$ *and* $\mathscr{Y}_2$ is defined by $Int(\mathscr{Y}_1, \mathscr{Y}_2) = \bigcup_{y_1 \in \mathscr{Y}_1, y_2 \in \mathscr{Y}_2} Int(y_1, y_2)$. The definition of interleaving is extended from two traces to n traces as follows: $Int(y_1, y_2, \ldots, y_n) = \bigcup_{y' \in Int(y_2, \ldots, y_n)} Int(y_1, y')$. For n sets $\mathscr{Y}_1, \mathscr{Y}_2, \ldots, \mathscr{Y}_n$ of observables, the interleaving $Int(\mathscr{Y}_1, \mathscr{Y}_2, \ldots, \mathscr{Y}_n)$ is defined analogously.

Although the interleaving defines all possible combinations of the sets of traces, they do not define the probability of their appearance. To reason about this, we define a scheduler that takes two sets of traces and probabilistically schedules their actions to form each possible trace in their interleaving.

Definition 5 (Scheduler) A *scheduler* $\mathscr{S}$ *on* $\mathscr{Y}_1$ *and* $\mathscr{Y}_2$ is a function that, given two traces $y_1 \in \mathscr{Y}_1$ and $y_2 \in \mathscr{Y}_2$, produces a probability distribution over all the possible interleaving $Int(y_1, y_2)$. We denote by $\mathscr{S}(y_1, y_2)[y]$ the conditional probability of having an interleaved trace y given y_1 and y_2.

We define a deterministic scheduler as one that produces the same output for any given two traces.

Definition 6 (Deterministic scheduler) A scheduler $\mathscr{S}$ is *deterministic* if for any two traces y_1 and y_2, there exists $y \in Int(y_1, y_2)$ such that $\mathscr{S}(y_1, y_2)[y] = 1$.

This provides the basis for composing channels in general, however this requires some delicacy since the interleaving of different traces may produce the same result. For example, given $y_1 = \tau.\overline{m}\langle s \rangle$ and $y_2 = \tau$ then one of the possible traces produced is $\tau.\tau.\overline{m}\langle s \rangle$. However, given $y_3 = \overline{m}\langle s \rangle$ and $y_4 = \tau.\tau$ then the same trace $\tau.\tau.\overline{m}\langle s \rangle$ could also be produced.

[1] Formally, we introduce $\mathscr{K}_{sep} = (\{sep\}, \{sep\}, (1))$ to consider the sequential execution of $\mathscr{K}_1$, $\mathscr{K}_{sep}$ and $\mathscr{K}_2$ in this order.

		observable			
		$\overline{m_1}\langle 0\rangle$	$\tau.\overline{m_1}\langle 0\rangle$	$\overline{m_1}\langle 1\rangle$	$\tau.\overline{m_1}\langle 1\rangle$
secret	0	0.5	0	0	0.5
	1	0	0.5	0.5	0

Table 1: Channel matrix C_1

		observable			
		$\overline{m_2}\langle 0\rangle$	$\tau.\overline{m_2}\langle 0\rangle$	$\overline{m_2}\langle 1\rangle$	$\tau.\overline{m_2}\langle 1\rangle$
secret	0	0	0.5	0.5	0
	1	0.5	0	0	0.5

Table 2: Channel matrix C_2

Let $p(y_1,y_2)$ be the joint probability that two component channels output two traces y_1 and y_2. Then the probability that $\mathscr{S}$ produces an interleaved trace y is given by: $p(y) = \sum_{y_1\in\mathscr{Y}_1,y_2\in\mathscr{Y}_2} p(y_1,y_2)\cdot \mathscr{S}(y_1,y_2)[y]$. By [20] we obtain $C_1[x_1,y_1]C_2[x_2,y_2] = p(y_1,y_2|x_1,x_2)$. Hence we can define scheduled composition of channels as follows.

Definition 7 (Scheduled composition of channels) The *scheduled composition of* two channels $\mathscr{K}_1 = (\mathscr{X}_1,\mathscr{Y}_2,C_1)$ and $\mathscr{K}_2 = (\mathscr{X}_2,\mathscr{Y}_2,C_2)$ *with respect to* a scheduler $\mathscr{S}$ is define as the channel $(\mathscr{X}_1\times\mathscr{X}_2, Int(\mathscr{Y}_1,\mathscr{Y}_2), C)$ where the matrix element for $x_1\in\mathscr{X}_1$, $x_2\in\mathscr{X}_2$ and $y\in Int(\mathscr{Y}_1,\mathscr{Y}_2)$ is given by: $C[(x_1,x_2),y] = \sum_{y_1\in\mathscr{Y}_1,y_2\in\mathscr{Y}_2} C_1[x_1,y_1]C_2[x_2,y_2].\mathscr{S}(y_1,y_2)[y]$.

We denote this scheduled composition by $Comp_{\mathscr{S}}(\mathscr{K}_1,\mathscr{K}_2)$. Note that the scheduled composition of n channels can be defined by adapting the scheduler $\mathscr{S}$ to operate over n traces in the obvious manner.

3.3 Examples of Scheduled Composition

This section presents some example channels and schedulers that illustrate the main results of this paper. For simplicity they shall all limit their secrets to the set $\mathscr{X}_B = \{0,1\}$, and their outputs to the set $\mathscr{Y}_m = \{\overline{m}\langle 0\rangle, \tau.\overline{m}\langle 0\rangle, \overline{m}\langle 1\rangle, \tau.\overline{m}\langle 1\rangle\}$ for a parameter m.

Consider the channel $\mathscr{K}_1 = (\mathscr{X}_B,\mathscr{Y}_{m_1},C_1)$ where C_1 is given by Table 1. This channel can be considered as one that half the time simply outputs the secret via $\overline{m_1}\langle s\rangle$ and half the time outputs the exclusive-or $\oplus$ of the secret with 1 as in $\tau.\overline{m_1}\langle s\oplus 1\rangle$, with the τ representing the calculation effort. Note that this channel leaks 100% of the information about the secret. Also consider the channel $\mathscr{K}_2 = (\mathscr{X}_B,\mathscr{Y}_{m_2},C_2)$ where C_2 is given by Table 2. This channel is similar to $\mathscr{K}_1$, except that the internal action τ is observable when disclosing the secret rather than its exclusive-or. Again this channel leaks all the secret information.

When combining channels the rôle of the scheduler is very significant with respect to the information leakage. This section defines three types of simple schedulers for illustrating the results here.

The simplest scheduler is one that outputs the first and second observable outputs concatenated, i.e. given y_1 and y_2 outputs $y_1.y_2$.

Definition 8 The *(left-first) deterministic sequential scheduler* $\mathscr{S}_{DS}$ is defined as follows: $\mathscr{S}_{DS}(y_1,y_2)[y]$ is 1 if $y = y_1.y_2$ and 0 otherwise where $y_1\in\mathscr{Y}_1$, $y_2\in\mathscr{Y}_2$ and $y\in\mathscr{Y}$.

Example 1 *The scheduled composition* $Comp_{\mathscr{S}_{DS}}(\mathscr{K}_1,\mathscr{K}_2)$ *w.r.t.* $\mathscr{S}_{DS}$ *has the same information leakage as the parallel composition* $\mathscr{K}_1\times\mathscr{K}_2$. *This can be shown since it follows from the definition of* $\mathscr{S}_{DS}$ *that, for each* $y\in\mathscr{Y}$, $\mathscr{S}_{DS}$ *uniquely identifies a pair* (y_1,y_2) *of outputs. For instance, let us consider the prior distribution* π *on* $\mathscr{X}_1\times\mathscr{X}_2$ *defined by* $(0.15, 0.20, 0.30, 0.35)$. *Then, for both of the composed channels, the mutual information is about* 1.926 *and the min-entropy leakage is about* 1.515.

Next is the *fair sequential scheduler* $\mathscr{S}_{FS}$ that fairly chooses between the first or second observable and produces that in its entirety before producing the other.

Definition 9 The *fair sequential scheduler* $\mathscr{S}_{FS}$ is defined by

$$\mathscr{S}_{FS}(y_1,y_2)[y] \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } y_1 = y_2 \wedge y = y_1.y_2 \\ 0.5 & \text{if } y_1 \neq y_2 \wedge (y = y_1.y_2 \vee y = y_2.y_1) \\ 0 & \text{otherwise.} \end{cases}$$

Similar to the deterministic sequential scheduler, the information leakage can be proven to be equal to that of the parallel composition of channels for this example.

Example 2 *The scheduled composition $\mathrm{Comp}_{\mathscr{S}_{FS}}(\mathscr{K}_1,\mathscr{K}_2)$ w.r.t. $\mathscr{S}_{FS}$ has the same information leakage as the parallel composition $\mathscr{K}_1 \times \mathscr{K}_2$. This can be shown similarly to Example 1.*

Note that the leakage preservation does *not* hold in general as illustrated in the following example.

Example 3 *Consider when $\mathscr{Y}_1 = \{\tau, \tau.\tau\}$ and $\mathscr{Y}_2 = \{\overline{m}\langle 0\rangle, \tau.\overline{m}\langle 0\rangle\}$. The observed output $\tau.\tau.\overline{m}\langle 0\rangle$ can arise from $\mathscr{S}(\tau, \tau.\overline{m}\langle 0\rangle)$ and $\mathscr{S}(\tau.\tau, \overline{m}\langle 0\rangle)$, where $\mathscr{S}$ can be $\mathscr{S}_{DS}$ or $\mathscr{S}_{FS}$. Thus, both the schedulers $\mathscr{S}_{DS}$ and $\mathscr{S}_{FS}$ may allow less information leakage than the parallel composition.*

The third example scheduler is the *fair interleaving scheduler* $\mathscr{S}_{FI}$ that evenly chooses the next action from the two observables.

Definition 10 The *fair interleaving scheduler* $\mathscr{S}_{FI}$ is recursively defined as

$$\mathscr{S}_{FI}(y_1,y_2)[y] \stackrel{\text{def}}{=} \begin{cases} 0.5\mathscr{S}_{FI}(y_1',y_2)[y'] & \text{if } y = \alpha.y' \wedge y_1 = \alpha.y_1' \wedge y_2 = \beta.y_2' \wedge \alpha \neq \beta \\ 0.5\mathscr{S}_{FI}(y_1,y_2')[y'] & \text{if } y = \beta.y' \wedge y_1 = \alpha.y_1' \wedge y_2 = \beta.y_2' \wedge \alpha \neq \beta \\ 0.5\mathscr{S}_{FI}(y_1',y_2)[y'] + 0.5\mathscr{S}_{FI}(y_1,y_2')[y'] & \text{if } y = \alpha.y' \wedge y_1 = \alpha.y_1' \wedge y_2 = \alpha.y_2' \\ 1 & \text{if } (y = y_1 \wedge y_2 = \emptyset) \vee (y = y_2 \wedge y_1 = \emptyset) \\ 0 & \text{otherwise.} \end{cases}$$

The fair interleaving scheduler $\mathscr{S}_{FI}$ turns out to often have impact on the leakage compared to the parallel composition of channels. This can occur in a variety of ways and shall be explored in detail later.

Example 4 *The scheduled composition $\mathrm{Comp}_{\mathscr{S}_F}(\mathscr{K}_1,\mathscr{K}_2)$ w.r.t. $\mathscr{S}_{FI}$ has less information leakage than the parallel composition $\mathscr{K}_1 \times \mathscr{K}_2$. This can be shown by considering when the output y is of the form $\tau.\overline{m_1}\langle 0\rangle.\overline{m_2}\langle 0\rangle$, which can arise from both $\mathscr{S}_{FI}(\tau.\overline{m_1}\langle 0\rangle, \overline{m_2}\langle 0\rangle)$ and $\mathscr{S}_{FI}(\overline{m_1}\langle 0\rangle, \tau.\overline{m_2}\langle 0\rangle)$. Since y does not uniquely identify the outputs y_1 and y_2, $\mathscr{S}_{FI}$ could allow less leakage than the parallel composition. For instance, for the prior $(0.15, 0.20, 0.30, 0.35)$, the mutual information of the scheduled composition w.r.t. $\mathscr{S}_{FI}$ is 1.695. This is less than those of the parallel composition and scheduled composition w.r.t. $\mathscr{S}_{DS}$ in Example 1 (both 1.926), thus the scheduler here alone is responsible for reducing the leakage.*

4 Information Leakage to Observers

4.1 Observers

Many kinds of capabilities of observing systems have been considered; e.g. an observer for strong bisimulation $\sim_s$ can recognise the internal action: $\tau.\overline{m}\langle v\rangle \not\sim_s \overline{m}\langle v\rangle$, while one for weak bisimulation $\sim_w$ cannot: $\tau.\overline{m}\langle v\rangle \sim_w \overline{m}\langle v\rangle$. To model different kinds of capabilities of observation, we define an *observer's views* $\mathscr{Z}$ as the set of values recognised by the observer. For example, $\tau.\overline{m}\langle v\rangle$ and $\overline{m}\langle v\rangle$ fall into two different views to an observer for strong bisimulation, but to the same view to an observer for weak bisimulation.

We formalise the notion of an observer using a matrix that defines relationships between observable outputs of systems and the observer's views. In particular, we allow for probabilistic accuracy in observation; that is the observer may not be perfectly accurate in identifying an output.

Definition 11 (Generalised observer) An *observer* $\mathscr{O}$ is defined as a triple $(\mathscr{Y},\mathscr{Z},Obs)$ consisting of a finite set $\mathscr{Y}$ of observables, a finite set $\mathscr{Z}$ of observer's views and an *observer matrix Obs* each of whose row represents a probability distribution; i.e., for all $y \in \mathscr{Y}$ we have $\sum_{z\in\mathscr{Z}} Obs[y,z] = 1$. Each matrix element $Obs[y,z]$ represents the probability that the observer has the view z when the actual output is y.

The observation matrix *Obs* describes the capability of the attacker to distinguish between traces. This capability of observation has been formalised as an equivalence relation between states of a system in prior work [6]. In fact, an equivalence relation $\sim$ between traces characterises a class of observers.

Definition 12 ($\sim$-observer) Given an equivalence relation $\sim$ on $\mathscr{Y}$, an observer $(\mathscr{Y},\mathscr{Z},Obs)$ is called a *$\sim$-observer* if, for all $y_1,y_2 \in \mathscr{Y}$, $y_1 \sim y_2$ is logically equivalent to $Obs[y_1,z] = Obs[y_2,z]$ for all $z \in Z$.

For instance, we can consider the $\sim_s$-observer for strong bisimulation $\sim_s$ and the $\sim_w$-observer for weak bisimulation $\sim_w$. Observe that $\sim_s$ is the identity relation on traces here. Further, note that for every observer $\mathscr{O}$, there exists an equivalence relation $\sim$ between traces such that $\mathscr{O}$ is a $\sim$-observer. This equivalence relation $\sim$ is defined by the following: $\sim \stackrel{\text{def}}{=} \{(y_1,y_2) \in \mathscr{Y}\times\mathscr{Y} \mid \text{ for all } z \in Z,\ Obs[y_1,z] = Obs[y_2,z]\}$. On the other hand, the observation matrix is *not* uniquely determined by the equivalence relation and therefore can express a wider range of observers' capabilities than the equivalence relation.

Among $\sim$-observers, we often consider observers that always have the same view on the same trace.

Definition 13 (Deterministic observer) We say that an observer $(\mathscr{Y},\mathscr{Z},Obs)$ is *deterministic* if each probability in *Obs* is either 0 or 1; i.e., for all $y \in \mathscr{Y}$, there exists a unique $z \in \mathscr{Z}$ such that $Obs[y,z] = 1$.

For any deterministic $\sim$-observer $(\mathscr{Y},\mathscr{Z},Obs)$ and any $y_1,y_2 \in \mathscr{Y}$, we have $y_1 \sim y_2$ iff, for all $z \in \mathscr{Z}$, we have $Obs[y_1,z] = Obs[y_2,z] \in \{0,1\}$. Then this observer always detects the equivalence class $[y]_\sim$ of the output y from any given view z. For this reason, when defining a deterministic $\sim$-observer, we typically take the set $\mathscr{Z}$ of views as the quotient set of $\mathscr{Y}$ by $\sim$, and for any $y \in \mathscr{Y}$ and $z \in \mathscr{Z}$, $Obs[y,z] = 1$ iff $z = [y]_\sim$. For example, consider the deterministic observers corresponding to $\sim_s$.

Example 5 (Deterministic $\sim_s$-observer) A deterministic $\sim_s$-observer $(\mathscr{Y},\mathscr{Z},Obs)$ satisfies the property that, for any distinct $y_1,y_2 \in \mathscr{Y}$, there exists a $z \in \mathscr{Z}$ such that either $Obs[y_1,z] = 0$ and $Obs[y_2,z] = 1$ or $Obs[y_1,z] = 1$ and $Obs[y_2,z] = 0$. Therefore this observer always detects the output y of the channel from any given view z. For this reason we call a deterministic $\sim_s$-observer a *perfect observer*.

Various kinds of bisimulations, or relations on observables, have been proposed and can be represented by various deterministic observers. Indeed, other kinds of relations can also be represented; consider an observer that cannot distinguish which source m_i a value is output upon. This can be formalised by using the equivalence relation $\sim_{ch}$ on traces that cannot distinguishes m_1 from m_2.

The last example observer here effectively ensures no leakage by seeing all outputs as the same:

Example 6 (Unit observer) An observer $\mathscr{O} = (\mathscr{Y},\mathscr{Z},Obs)$ is called a *unit observer* if $\mathscr{Z}$ is a singleton. It has the same view regardless of the outputs of the channel, thus can detect no leakage of the channel.

4.2 Observed Information Leakage

The amount of observed information leakage depends on the capability of the observer. To quantify this we introduce the notion of *observed information leakage*.

Definition 14 (Observed information leakage) Let $\mathscr{K} = (\mathscr{X},\mathscr{Y},C)$ be a channel and $\mathscr{O} = (\mathscr{Y},\mathscr{Z},Obs)$ be an observer. For each leakage measure $L \in \{\mathscr{I},\mathscr{L}\}$ and any prior π on $\mathscr{X}$, we define *observed information leakage* by $L_{\mathscr{O}}(\pi,\mathscr{K}) = L(\pi,\mathscr{K}\cdot\mathscr{O})$ where $\mathscr{K}\cdot\mathscr{O} = (\mathscr{X},\mathscr{Z},C\cdot Obs)$ is the cascade composition [18] of $\mathscr{K}$ and $\mathscr{O}$. Similarly, for each $L \in \{\mathscr{SC},\mathscr{MC}\}$, we define $L_{\mathscr{O}}(\mathscr{K}) = L(\mathscr{K}\cdot\mathscr{O})$.

We present properties of observed information leakage as follows. The first remark is that, for each equivalence relation $\sim$ on traces, all deterministic $\sim$-observers give the same observed leakage values.

Proposition 1 *Let π be any prior on $\mathscr{X}$ and $\mathscr{K} = (\mathscr{X}, \mathscr{Y}, C)$ be any channel. For any equivalence relation $\sim$ on $\mathscr{Y}$ and any two deterministic $\sim$-observers $\mathscr{O}_1$, $\mathscr{O}_2$, we have $L_{\mathscr{O}_1}(\pi, \mathscr{K}) = L_{\mathscr{O}_2}(\pi, \mathscr{K})$ for $L \in \{\mathscr{I}, \mathscr{L}\}$ and $L_{\mathscr{O}_1}(\mathscr{K}) = L_{\mathscr{O}_2}(\mathscr{K})$ for $L \in \{\mathscr{SC}, \mathscr{MC}\}$.*

The following states that the deterministic $\sim_s$-observers and unit observers respectively have the maximum and minimum capabilities of distinguishing traces. That is, the deterministic $\sim_s$-observer can detect every behaviour of the channel accurately and does not alter the leakage of the channel in any manner, while the unit observers cannot detect any leakage of the channel.

Proposition 2 *For each $L \in \{\mathscr{I}, \mathscr{L}\}$, $0 \leq L_{\mathscr{O}}(\pi, \mathscr{K}) \leq L(\pi, \mathscr{K})$. For each $L \in \{\mathscr{SC}, \mathscr{MC}\}$, $0 \leq L_{\mathscr{O}}(\mathscr{K}) \leq L(\mathscr{K})$. In these inequations, the left equalities hold when $\mathscr{O}$ is a unit observer, and the right ones hold when $\mathscr{O}$ is a deterministic $\sim_s$-observer.*

Next we compare the capabilities of generalised observers. Recall the composition-refinement relation $\sqsubseteq_\circ$ on channels [3, 24]: A channel $\mathscr{K}_1$ is *composition-refined* by another $\mathscr{K}_2$, written as $\mathscr{K}_1 \sqsubseteq_\circ \mathscr{K}_2$, iff there exists a channel $\mathscr{K}'$ such that $\mathscr{K}_1 = \mathscr{K}_2 \cdot \mathscr{K}'$. Since the generalised observers are also channels, we can consider this ordering $\sqsubseteq_\circ$ on observers. For example, the unit observer is composition-refined by $\sim_w$-observers, and the deterministic $\sim_w$-observer is by the deterministic $\sim_s$-observer. For another example, any probabilistic $\sim_a$-observer is composition-refined by the deterministic $\sim_a$-observer:

Proposition 3 *Given any equivalence relation $\sim_a$ on $\mathscr{Y}$ let $\mathscr{O}_1 = (\mathscr{Y}, \mathscr{Z}, Obs_1)$ and $\mathscr{O}_2 = (\mathscr{Y}, \mathscr{Z}, Obs_2)$ be two $\sim_a$-observers. If $\mathscr{O}_2$ is deterministic then $\mathscr{O}_1 \sqsubseteq_\circ \mathscr{O}_2$.*

The composition-refined observer will observe less information leakage.

Theorem 4 *Let $\mathscr{O}_1$ and $\mathscr{O}_2$ be two observers such that $\mathscr{O}_1 \sqsubseteq_\circ \mathscr{O}_2$. Then, for any prior π and any channel $\mathscr{K}$, we have $L_{\mathscr{O}_1}(\pi, \mathscr{K}) \leq L_{\mathscr{O}_2}(\pi, \mathscr{K})$ for $L \in \{\mathscr{I}, \mathscr{L}\}$ and $L_{\mathscr{O}_1}(\mathscr{K}) \leq L_{\mathscr{O}_2}(\mathscr{K})$ for $L \in \{\mathscr{SC}, \mathscr{MC}\}$.*

These results imply that no probabilistic $\sim$-observer detect more leakage than deterministic ones.

4.3 Examples of Deterministic Observers

Theorem 4 implies that the deterministic $\sim_s$-observer does not observe less information leakage than the deterministic $\sim_w$-observer.

Example 7 *Let us consider the scheduled compositions in Examples 1 and 2 in Section 3.3. Both the composed channels leak all secrets without considering observers; i.e., they do so in the presence of $\sim_s$-observer. On the other hand, they leak no secrets to a weakly-bisimilar observer. For example, for each $i \in \{1,2\}$, we define the deterministic $\sim_w$-observer $\mathscr{O}_i$ as $(\{\overline{m_i}\langle 0\rangle, \tau.\overline{m_i}\langle 0\rangle, \overline{m_i}\langle 1\rangle, \tau.\overline{m_i}\langle 1\rangle\}, \{[\overline{m_i}\langle 0\rangle]_{\sim_w}, [\overline{m_i}\langle 1\rangle]_{\sim_w}\}, Obs)$ where Obs is the matrix given in Table 3. Applying the $\sim_w$-observer $\mathscr{O}_i$ to both $\mathscr{K}_1$ and $\mathscr{K}_2$ yields the same matrix presented in Table 4. Then both channels leak no information to the $\sim_w$-observer. Therefore, the deterministic $\sim_s$-observer observes more information leakage than the deterministic $\sim_w$-observer also in this example.*

The scheduled composition can also leak more information than the parallel composition (and even than each component channel) in the presence of imperfect observers.

Example 8 (Observer dependent) *Consider the scheduled composition of the channels $\mathscr{K}_1$ and $\mathscr{K}_2$ w.r.t. the fair interleaving scheduler $\mathscr{S}_{FI}$. By Example 4, the leakage of the scheduled composition w.r.t. $\mathscr{S}_{FI}$ is less than that of the parallel composition in the presence of the deterministic $\sim_s$-observer.*

However, the leakage of the scheduled composition is more than that of the parallel composition (and even than that of each component channel) when the $\sim_w$-observer $\mathscr{O}$ is being considered; e.g., $\mathscr{L}_{\mathscr{O}}(\pi, Comp_{\mathscr{S}_{FI}}(\mathscr{K}_1, \mathscr{K}_2)) = 0.215 > 0 = \mathscr{L}_{\mathscr{O}}(\pi, \mathscr{K}_1 \times \mathscr{K}_2) = \mathscr{L}_{\mathscr{O}}(\pi, \mathscr{K}_1)$ for $\pi = (0.15, 0.20, 0.30, 0.35)$.

		view	
		$[\overline{m_i}\langle 0\rangle]_{\sim_w}$	$[\overline{m_i}\langle 1\rangle]_{\sim_w}$
output	$\overline{m_i}\langle 0\rangle$ or $\tau.\overline{m_i}\langle 0\rangle$	1	0
	$\overline{m_i}\langle 1\rangle$ or $\tau.\overline{m_i}\langle 1\rangle$	0	1

Table 3: Observer matrix *Obs*

		view	
		$[\overline{m_i}\langle 0\rangle]_{\sim_w}$	$[\overline{m_i}\langle 1\rangle]_{\sim_w}$
secret	0	0.5	0.5
	1	0.5	0.5

Table 4: Composed matrix $C_i \cdot Obs$

4.4 Example of Probabilistic Observers

The notion of deterministic $\sim$-observers is useful to model various observers, but they may not cover all realistic settings. For example, when the internal action τ represents time to perform internal computation, observers may recognise it only probabilistically, for instance with probability 0.7. Then such *probabilistic observers* cannot be modeled as deterministic observers but as generalised observers, which quantify the capabilities of probabilistic observation. As far as we know, no previous work on quantitative information flow analyses have considered probabilistic observers.

Example 9 *Consider a probabilistic observer $\mathscr{O}$ that can recognise a single internal action τ only probabilistically but two or more consecutive τ's with probability* 1*. For instance, $\mathscr{O}$ recognises the trace $(\tau.\overline{m_i}\langle 0\rangle.\overline{m_i}\langle 1\rangle)$ correctly with probability* 0.7 *and confuses it with either $(\overline{m_i}\langle 0\rangle.\overline{m_i}\langle 1\rangle)$, $(\overline{m_i}\langle 0\rangle.\tau.\overline{m_i}\langle 1\rangle)$ or $(\overline{m_i}\langle 0\rangle.\overline{m_i}\langle 1\rangle.\tau)$ each with probability* 0.1*. Consider the schedule-composed channel $Comp_{\mathscr{S}_{FI}}(\mathscr{K}_1,\mathscr{K}_2)$ from Example 4. The observed mutual information is* 0.783 *under the probabilistic observer $\mathscr{O}$, which is between* 0.090 *and* 1.695 *as observed under the deterministic $\sim_w$-observer and $\sim_s$-observer.*

5 Relationships between Scheduling and Observation

This section generalises the previous examples to show three kinds of results. First, we identify conditions on component channels under which leakage cannot be effected by the scheduled composition. Second, we show that scheduled composition can leak more or less information than the parallel composition, including results on the bounds of the information leaked. Third, we present an algorithm for finding a scheduler that minimises the min-entropy leakage/min-capacity under any observer

5.1 Information Leakage Independent of Scheduling

This section presents results for determining when the leakage is independent of the scheduler. Regardless of the scheduler and observer, the leakage of the scheduled composition is equivalent to that of the parallel composition under certain conditions on component channels that are detailed below.

Theorem 5 *Let $\mathscr{K}_1 = (\mathscr{X}_1,\mathscr{Y}_1,C_1)$ and $\mathscr{K}_2 = (\mathscr{X}_2,\mathscr{Y}_2,C_2)$ be channels. Assume that, for any $y_1,y_1' \in \mathscr{Y}_1$ and $y_2,y_2' \in \mathscr{Y}_2$, if $Int(y_1,y_2) \cap Int(y_1',y_2') \neq \emptyset$ then $y_1 = y_1'$ and $y_2 = y_2'$. Then, for every scheduler $\mathscr{S}$ and observer $\mathscr{O}$, the leakage of the scheduled composition is the same as that of the parallel composition.*

By adding a stronger requirement to Theorem 5, we obtain the following corollary.

Corollary 6 *Let $\mathscr{K}_1 = (\mathscr{X}_1,\mathscr{Y}_1,C_1)$ and $\mathscr{K}_2 = (\mathscr{X}_2,\mathscr{Y}_2,C_2)$ be channels. Assume that, for all $(y_1,y_2) \in \mathscr{Y}_1 \times \mathscr{Y}_2$, $\alpha \in y_1$ and $\beta \in y_2$, we have $\alpha \neq \beta$. Then, for every scheduler $\mathscr{S}$ and observer $\mathscr{O}$, the leakage of the scheduled composition is the same as that of the parallel composition.*

5.2 Schedulers for Altering Information Leakage

This section considers when schedulers can alter the leakage of a scheduled composition. This is distinct from prior results where it has been shown that the composition cannot leak more information than the component channels [5, 19, 20], since here more information can be leaked to imperfect observers.

In general scheduled composition can yield more or less leakage than the individual component channels or their parallel composition. This is illustrated by Example 8. Unfortunately heuristics for determining when more information is leaked end up being rather complicated and dependent on many relations between traces, interleavings, equivalences, and then subject to generalities about both schedulers and observers. Ultimately it is easier to show by examples that, for some channels, prior, and $\sim$-observer, there is a scheduler by which the scheduled composition leaks strictly more information than the parallel composition. Since this clearly holds by example, we consider a class of schedulers under which the scheduled composition does not leak more information than the parallel composition.

To define this we extend an equivalence relation $\sim$ on traces to probability distributions of traces: We say that two distributions D and D' on a set $\mathscr{Y}$ are $\sim$-*indistinguishable* (written as $D \sim D'$) if the deterministic $\sim$-observer cannot distinguish D from D' at all, i.e., for all equivalence class $t \in \mathscr{Y}/\sim$, we have $\sum_{y\in t} D[y] = \sum_{y\in t} D'[y]$. Using $\sim$-indistinguishability we define a scheduler that does not leak any behaviour of the system that the $\sim$-observer cannot detect.

Definition 15 Let $\sim$ be an equivalence relation on $\mathscr{Y}_1 \cup \mathscr{Y}_2 \cup \mathit{Int}(\mathscr{Y}_1, \mathscr{Y}_2)$. A scheduler $\mathscr{S}$ on $\mathscr{Y}_1$ and $\mathscr{Y}_2$ is a $\sim$-*blind scheduler* when, for any two pairs $(y_1,y_2),(y_1',y_2') \in \mathscr{Y}_1 \times \mathscr{Y}_2$, we have $y_1 \sim y_1'$ and $y_2 \sim y_2'$ iff we have $\mathscr{S}(y_1,y_2) \sim \mathscr{S}(y_1',y_2')$.

For instance, the deterministic sequential scheduler $\mathscr{S}_{DS}$ and the fair sequential scheduler $\mathscr{S}_{FS}$ are $\sim_w$-blind while the fair interleaving scheduler $\mathscr{S}_{FI}$ is not. Note that $\sim$-blind schedulers do not leak any behaviour that would not be visible to the deterministic $\sim$-observers. Thus they do not gain more information from the scheduled composition w.r.t. $\sim$ than the parallel composition.

Theorem 7 *Let π be a prior, $\mathscr{K}_1$ and $\mathscr{K}_2$ be two channels, $\mathscr{O}$ be a deterministic $\sim$-observer, and S be a $\sim$-blind scheduler. For each $L \in \{\mathscr{I},\mathscr{L}\}$ we have $L_{\mathscr{O}}(\pi, \mathit{Comp}_{\mathscr{S}}(\mathscr{K}_1,\mathscr{K}_2)) \leq L_{\mathscr{O}}(\pi, \mathscr{K}_1 \times \mathscr{K}_2)$. For each $L \in \{\mathscr{SC},\mathscr{MC}\}$ we have $L_{\mathscr{O}}(\mathit{Comp}_{\mathscr{S}}(\mathscr{K}_1,\mathscr{K}_2)) \leq L_{\mathscr{O}}(\mathscr{K}_1 \times \mathscr{K}_2)$. When $\mathscr{S}$ is also deterministic, the leakage relations become equalities.*

For instance, since $\mathscr{S}_{DS}$ and $\mathscr{S}_{FS}$ are $\sim_w$-blind schedulers, the deterministic $\sim_w$-observers do not gain more information from the scheduled composition w.r.t. $\sim_w$ than the parallel composition. In fact, they have the same leakage in Example 7.

The following result is about a heuristic for when leakage can be changed by the properties of the scheduler. This is presented here to clarify the properties.

Theorem 8 *Let $\mathscr{K}_1 = (\mathscr{X}_1,\mathscr{Y}_1,C_1)$ and $\mathscr{K}_2 = (\mathscr{X}_2,\mathscr{Y}_2,C_2)$ be two channels. Assume that there exist $y_1,y_1' \in \mathscr{Y}_1$ and $y_2,y_2' \in \mathscr{Y}_2$ such that $\mathit{Int}(y_1,y_2) \cap \mathit{Int}(y_1',y_2') \neq \emptyset$. Then it is possible for the scheduled-composition of $\mathscr{K}_1$ and $\mathscr{K}_2$ to alter the mutual information and min-entropy leakage for some prior.*

5.3 Schedulers for Minimising Information Leakage

This section presents results for finding a scheduler that minimises the min-entropy leakage and min-capacity in the presence of any observer.

Theorem 9 *Given any prior π, two channels $\mathscr{K}_1$, $\mathscr{K}_2$ and any observer $\mathscr{O}$, there is an algorithm that computes a scheduler $\mathscr{S}$ that minimises the observed min-entropy leakage $\mathscr{L}_{\mathscr{O}}(\pi, \mathit{Comp}_{\mathscr{S}}(\mathscr{K}_1,\mathscr{K}_2))$ of the scheduled composition.*

Proof: To find a scheduler $\mathscr{S}$ that minimises the observed min-entropy leakage $\mathscr{L}_{\mathscr{O}}(\pi, Comp_{\mathscr{S}}(\mathscr{K}_1,\mathscr{K}_2))$, it is sufficient to find $\mathscr{S}$ that minimises the observed posterior vulnerability $V(\pi, Comp_{\mathscr{S}}(\mathscr{K}_1,\mathscr{K}_2)\cdot\mathscr{O})$.

For $(x_1,x_2)\in\mathscr{X}_1\times\mathscr{X}_2$ and $(y_1,y_2)\in\mathscr{Y}_1\times\mathscr{Y}_2$, let $p(x_1,x_2,y_1,y_2)=\pi[x_1,x_2](C_1\times C_2)[(x_1.x_2),(y_1,y_2)]$. For each $z\in\mathscr{Z}$ let $v_z = \max_{(x_1,x_2)\in\mathscr{X}_1\times\mathscr{X}_2}\sum_{y_1,y_2,y} p(x_1,x_2,y_1,y_2)\mathscr{S}(y_1,y_2)[y]Obs[y,z]$ where (y_1,y_2) and y range over $\mathscr{Y}_1\times\mathscr{Y}_2$ and $Int(\mathscr{Y}_1,\mathscr{Y}_2)$ respectively. Let $Pos(y_1,y_2)[y]$ be the $(|\mathscr{Y}_1|\times|\mathscr{Y}_2|, |Int(\mathscr{Y}_1,\mathscr{Y}_2)|)$-matrix defined by the following: $Pos(y_1,y_2)[y]=1$ if y can be obtained by interleaving y_1 and y_2, and $Pos(y_1,y_2)[y]=0$ otherwise.

To find a scheduler matrix $\mathscr{S}$ that minimises the observed posterior vulnerability, it suffices to solve the linear program that minimises $\sum_{z\in\mathscr{Z}} v_z$, subject to

- for each $(x_1,x_2,z)\in\mathscr{X}_1\times\mathscr{X}_2\times\mathscr{Z}$, $\sum_{y_1,y_2,y} p(x_1,x_2,y_1,y_2)\mathscr{S}(y_1,y_2)[y]Obs[y,z]\leq v_z$
- for each $(y_1,y_2)\in\mathscr{Y}_1\times\mathscr{Y}_2$, $\sum_y Pos(y_1,y_2)[y]\mathscr{S}(y_1,y_2)[y]=1$.

Note that the second constraint means that each row of the scheduler matrix $\mathscr{S}$ must sum to 1. In this linear program, the scheduler matrix element $\mathscr{S}(y_1,y_2)[y]$ for each $(y_1,y_2)\in\mathscr{Y}_1\times\mathscr{Y}_2$ and $y\in Int(\mathscr{Y}_1,\mathscr{Y}_2)$ and v_z for each $z\in\mathscr{Z}$ are variables. We can solve this problem using the simplex method or interior point method. (In practice, we can efficiently solve it using a linear programming solver such as lp_solve [2].) Hence we obtain a scheduler matrix $\mathscr{S}$ that minimises $\sum_{z\in\mathscr{Z}} v_z$. □

In the above linear program the number of variables is $|\mathscr{Y}_1|\times|\mathscr{Y}_2|\times|Int(\mathscr{Y}_1,\mathscr{Y}_2)|+|\mathscr{Z}|$, and the number of constraints is $|\mathscr{X}_1|\times|\mathscr{X}_2|\times|\mathscr{Z}|+|\mathscr{Y}_1|\times|\mathscr{Y}_2|$. Since the number of interleaved traces grows exponentially in the number of traces, the time to compute a minimising scheduler is exponential in the number of component traces. When the observer $\mathscr{O}$ is imperfect enough for $|\mathscr{Z}|$ to be very small, then the computation time improves significantly in practice. On the other hand, when the number of traces is very large, we may heuristically obtain a scheduler with less leakage by results in the previous section.

To obtain a scheduler that minimises the worst-case leakage value, it suffices to consider a scheduler that minimises the min-capacity.

Corollary 10 *Given two channels $\mathscr{K}_1$, $\mathscr{K}_2$ and any observer $\mathscr{O}$, there is an algorithm that computes a scheduler $\mathscr{S}$ that minimises the observed min-capacity of the scheduled composition.*

These two results give the minimum amount of leakage that is possible for any scheduling.

Example 10 *Consider the channels $\mathscr{K}_1,\mathscr{K}_2$ defined in Section 3.3. By Theorem 9, the minimum observed min-entropy leakage w.r.t. the prior* $(0.15, 0.20, 0.30, 0.35)$ *is* 1.237 *under the deterministic $\sim_s$-observer, and* 0.801 *under the probabilistic observer defined in Example 9. By Corollary 10, the minimum observed min-capacity is* 1.585 *under the deterministic $\sim_s$-observer, and* 1.138 *under the probabilistic observer.*

Since the channel capacity will not exceed the min-capacity [26], the minimum observed min-capacity obtained by the above algorithm gives an upper bound on the minimum channel capacity.

6 Case Studies

6.1 Sender Anonymity

In e-voting *sender anonymity* can be summarised as the issue of collecting votes from a number of voters and being able to expose the aggregate vote information while revealing as little as possible about how each voter voted. This can be solved by a general application of a mix network [13] where all the votes are sent via mixing systems that output the votes in a manner that should not reveal how each voter voted.

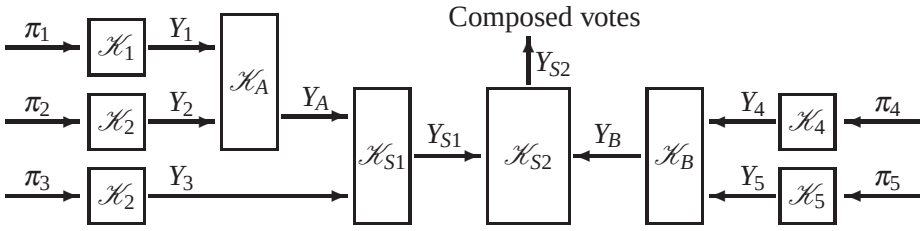

Figure 2: Structure of composed channels for voters

This can be represented here by each voter being an information-theoretic channel that outputs their vote. For example, consider a simple voting in which possible votes are 0 and 1 and each voter outputs the chosen vote via $\overline{m}\langle 0\rangle$ or $\overline{m}\langle 1\rangle$, respectively. Then each voter (indexed by i) can be represented by the channel $\mathscr{K}_i = (\{0,1\}, \{\overline{m}\langle 0\rangle, \overline{m}\langle 1\rangle\}, C_i)$ where $C_i[k, \overline{m}\langle k\rangle] = 1$ for $k \in \{0,1\}$ and each voter has a prior π_i on $\{0,1\}$. Observe that each such voter channel alone fully reveals the prior for the channel.

The scheduled composition of the voters represents the mix network with the schedulers representing the mixing algorithm and thus providing the ability to reason over their effect on information leakage. Consider the following problem with five voters $\mathscr{K}_1$ to $\mathscr{K}_5$. As illustrated in Figure 2, the ballot of each voter is sent via intermediate severs (schedulers) $\mathscr{K}_A$, $\mathscr{K}_B$, $\mathscr{K}_{S1}$ that mix the order of ballots. The final system $\mathscr{K}_{S2}$ combines $\mathscr{K}_{S1}$ and $\mathscr{K}_B$ to output all the votes according to some mixing.

Using the deterministic sequential scheduler $\mathscr{S}_{DS}$ for all compositions reveals all information on how each voter voted. That is, the leakage is considered to be 5-bits (as each vote is 0 or 1). On the other hand, using the fair sequential scheduler $\mathscr{S}_{FS}$ for all compositions leaks less information than $\mathscr{S}_{DS}$. When π is uniform and $\mathscr{K}$ is the composed channel in Figure 2 with the appropriate scheduling, we obtain $\mathscr{L}(\pi, \mathscr{K}) = 3.426$ and $\mathscr{I}(\pi, \mathscr{K}) = 2.836$. Observe that here the third voter's output can only appear in the 1st, 3rd, or 5th position in the final trace. This is repaired by using the fair interleaving scheduler $\mathscr{S}_{FI}$ for all compositions that leaks even less information: $\mathscr{L}(\pi, \mathscr{K}) = 2.901$ and $\mathscr{I}(\pi, \mathscr{K}) = 2.251$.

A more interesting case is when different compositions use different schedulers. Since the votes do not contain any information about the system they came from, let alone voter. Using the fair sequential scheduler for $\mathscr{K}_A$ and $\mathscr{K}_B$, and the fair interleaving scheduler for $\mathscr{K}_{S2}$, along with a specially constructed scheduler for $\mathscr{K}_{S1}$ can reduce the information leakage to a minimum. Then the min-entropy leakage is 2.824 and the mutual information is 2.234. Note that when there is only one scheduler that receives all 5 ballots, the minimum min-capacity of the composed system (over all possible schedulers) is 2.585.

The example can be extended further by adding τ steps before votes to indicate time taken for some parts of the process. For a simple example, consider when voters 1 and 2 have a τ step before their vote to represent the time taken, e.g. as indicative of voting order, or the time taken for the extra mixing step. In the presence of all fair interleaving schedulers, the observed min-entropy leakage and the mutual information are respectively 3.441 and 2.785 under the perfect observer. However, these shift to 3.381 and 2.597, respectively, under the deterministic $\sim_w$-observer.

6.2 Side-Channel Attacks

Consider the small program shown in Figure 3, where an observable action is repeated in a loop. This program captures, for instance, some aspects of decryption algorithms of certain cryptographic schemes, such as RSA. Intuitively, `X[]` is the binary array representing a 3-bit secret (e.g. `011`), which corresponds to secret decryption keys. The timing of the algorithm's operation reveals which bit of the secret key is 1, since the observable-operation $\overline{m}\langle 1\rangle$ can be detected, perhaps as power consumption, response time, or some other side-effect of the algorithm [22]. We denote by $\mathscr{K}$ the channel defined by this program.

```
for(i = 0; i < 3; i++) {
   tau;
   if(X[i] = 1) {
     m<1>; //observable-operation
   }
}
```

Figure 3: Decryption algorithm

		view		
		τ	$\overline{m}\langle 1\rangle$	$\emptyset$
output	τ	0.8	0.1	0.1
	$\overline{m}\langle 1\rangle$	0.05	0.9	0.05

Figure 4: Probabilistic observer matrix

Consider composition of $\mathscr{K}$ with itself, e.g., when applying the algorithm to different parts of the message in parallel. Clearly if the parallel composition is taken then both instances of $\mathscr{K}$ will leak all their information about the key. On the other hand, the scheduled composition may have less leakage.

We first consider the case each instance of the component channel $\mathscr{K}$ receives a different secret bit string independently drawn from the uniform prior. This captures the situation in which each decryption operation uses different secret keys. When the fair interleaving scheduler mixes the two traces, the min-entropy leakage and the mutual information are respectively 4.257 and 3.547 in the presence of the perfect observer, and 2.807 and 2.333 in the presence of the deterministic $\sim_w$-observer.

Next we consider the case where both instances of $\mathscr{K}$ share the same secret key which has been drawn from the uniform prior. When the fair interleaving scheduler mixes the two traces, the min-entropy leakage and the mutual information are respectively 3.000 and 3.000 (all 3 bits of the secret key are leaked) under the perfect observer, and 2.000 and 1.811 under the deterministic $\sim_w$-observer.

More interesting is to consider the case where the observer is only able to detect approximate information through the side-channel. Consider the observer $\mathscr{O}$ that only probabilistically observes actions according to the matrix in Figure 4. Here $\emptyset$ indicates that nothing is detected by the attacker not even a τ. For example, applying this observer to the trace $\tau.\tau.\tau$ may yield $\tau.\tau$ when one τ is not observed (represented $\emptyset$ in the matrix). Such an observer is less effective even when applied to the parallel composition of channels. However, this applies even further when applied to any scheduled composition since the loss of information through poor detection cannot even be limited to one channel or the other. Thus, a trace of length 5, even from a leaky scheduler such as the (left-first) sequential scheduler, would leak less than the parallel composition (since it would be clear which composite channel had been poorly observed).

For instance, let us consider the case each instance of $\mathscr{K}$ independently receives a secret from the uniform prior and the fair interleaving scheduler is used. Then the min-entropy leakage and the mutual information are respectively 3.306 and 1.454 under this probabilistic observer. If we consider the case both instances of $\mathscr{K}$ shares the same secret, then the leakage values are respectively 2.556 and 1.924.

7 Related Work

Regarding schedulers there are a variety of studies on relationships between schedulers and information leakage [11, 4]. In [10] the authors consider a *task-scheduler* that is similar to our schedulers, albeit restricted to the form of our deterministic scheduler. The schedulers in this paper are also similar to the *admissible schedulers* of [4]. Both are defined to depend only upon the observable outputs, that is the traces they schedule. This avoids the possibility of leakage via the scheduler being aware of the intended secret directly and so leaking information. Differently to admissible schedulers, here the scheduler can be probabilistic, which is similar in concept to the probabilistically defined (deterministic) schedulers of [27], although they explore scheduling and determinism of Markov Chains and not information leakage.

Most work on schedulers has focused on preventing any leakage at all, indeed the problem is typically defined to prevent any high/secret information leaking. This in turn sets extremely high requirements upon the scheduler, and so proves to be difficult to achieve, or even impossible. Here we take an approach to scheduling that allows for probabilistic schedulers and so reasoning about the quantitative information

leakage, rather than total leakage. Thus we permit schedulers that can be daemonic or angelic, as well as many in between that may closer resemble the behaviour of real world systems.

Regarding observers there is little prior work in quantitative information flow and quantifying the capability of the observer. [6] has some similarity where they formalise an equivalence of system states similar in style to the deterministic $\sim$-observers here. However, this does not model observers as part of information-theoretic channels, hence does not allow the probabilistic behaviour of observers.

8 Conclusions and Future Work

We have introduced the notion of the scheduled composition of channels and generalised the capabilities of the observers to reason about more systems. Then we have presented theories that can be used as heuristics to detect when scheduled composition may have an effect on the information leakage. This determines when scheduled composition is a potential risk/benefit to a scheduler-dependent system. Scheduling can both leak more information, or less information to an observer depending on many factors, while some leakage bounds can be obtained for schedule-composed channels. Further, we have shown an algorithm for finding a scheduler that minimises the leakage of the scheduled composition.

The work here provides a foundation for continuing research into concurrent behavior, including interactive systems. Here we have limited the systems to finite sets of secrets and observables since this aligns with the discrete version of leakage calculations. By shifting to continuous domains we can investigate some systems with infinite secrets or observables. Similarly the schedulers here assume finite traces and are typically defined over the entire possible traces. However, many do not require this, and can be defined only upon the next action in the trace. This allows for alternate definitions without changing the results, and easier applicability to infinite settings.

References

[1] `http://www.cs.bham.ac.uk/research/projects/infotools/leakiest/ext/`.

[2] *lp_solve version 5.5*. `http://lpsolve.sourceforge.net/`.

[3] Mário S. Alvim, Konstantinos Chatzikokolakis, Catuscia Palamidessi & Geoffrey Smith (2012): *Measuring Information Leakage Using Generalized Gain Functions*. In: *Proc. of CSF*, IEEE, pp. 265–279, doi:`10.1109/CSF.2012.26`.

[4] Miguel Andres, E., Catuscia Palamidessi, Ana Sokolova & Peter Van Rossum (2011): *Information Hiding in Probabilistic Concurrent Systems. Theor. Comp. Sci.* 412(28), pp. 3072–3089, doi:`10.1016/j.tcs.2011.02.045`.

[5] Gilles Barthe & Boris Köpf (2011): *Information-theoretic Bounds for Differentially Private Mechanisms*. In: *Proc. of CSF*, IEEE, pp. 191–204, doi:`10.1109/CSF.2011.20`.

[6] Fabrizio Biondi, Axel Legay, Pasquale Malacaria & Andrzej Wasowski (2013): *Quantifying Information Leakage of Randomized Protocols*. In: *Proc. of VMCAI*, pp. 68–87, doi:`10.1007/978-3-642-35873-9_7`.

[7] Michele Boreale (2009): *Quantifying information leakage in process calculi. Inf. Comput.* 207(6), pp. 699–725, doi:`10.1016/j.ic.2008.12.007`.

[8] Michele Boreale, Francesca Pampaloni & Michela Paolini (2011): *Asymptotic Information Leakage under One-Try Attacks*. In: *Proc. of FOSSACS*, pp. 396–410, doi:`10.1007/978-3-642-19805-2_27`.

[9] Christelle Braun, Konstantinos Chatzikokolakis & Catuscia Palamidessi (2009): *Quantitative Notions of Leakage for One-try Attacks*. In: *Proc. of MFPS*, *ENTCS* 249, Elsevier, pp. 75–91, doi:`10.1016/j.entcs.2009.07.085`.

[10] Ran Canetti, Ling Cheung, Dilsun Kirli Kaynar, Moses Liskov, Nancy A. Lynch, Olivier Pereira & Roberto Segala (2008): *Analyzing Security Protocols Using Time-Bounded Task-PIOAs*. *Discrete Event Dynamic Systems* 18, pp. 111–159, doi:10.1007/s10626-007-0032-1.

[11] Konstantinos Chatzikokolakis & Catuscia Palamidessi (2007): *Making random choices invisible to the scheduler*. In: *Proc. of CONCUR'07*, Springer, pp. 42–58, doi:10.1016/j.ic.2009.06.006.

[12] Konstantinos Chatzikokolakis, Catuscia Palamidessi & Prakash Panangaden (2008): *Anonymity Protocols as Noisy Channels*. *Inf. Comput.* 206(2–4), pp. 378–401, doi:10.1016/j.ic.2007.07.003.

[13] David Chaum (1981): *Untraceable electronic mail, return addresses, and digital pseudonyms*. *Commun. ACM* 24(2), pp. 84–90, doi:10.1145/358549.358563.

[14] Tom Chothia, Yusuke Kawamoto & Chris Novakovic (2013): *A Tool for Estimating Information Leakage*. In: *Proc. of CAV'13*, doi:10.1007/978-3-642-39799-8_47.

[15] Tom Chothia, Yusuke Kawamoto & Chris Novakovic (2014): *LeakWatch: Estimating Information Leakage from Java Programs*. In: *Proc. of ESORICS'14*, pp. 219–236, doi:10.1007/978-3-319-11212-1_13.

[16] Tom Chothia, Yusuke Kawamoto, Chris Novakovic & David Parker (2013): *Probabilistic Point-to-Point Information Leakage*. In: *Proc. of CSF*, IEEE, pp. 193–205, doi:10.1109/CSF.2013.20.

[17] David Clark, Sebastian Hunt & Pasquale Malacaria (2001): *Quantitative Analysis of the Leakage of Confidential Data*. In: *Proc. of QAPL'01*, *ENTCS* 59 (3), Elsevier, pp. 238–251, doi:10.1016/S1571-0661(04)00290-7.

[18] Barbara Espinoza & Geoffrey Smith (2011): *Min-Entropy Leakage of Channels in Cascade*. In: *Proc. of FAST*, *LNCS* 7140, Springer, pp. 70–84, doi:10.1007/978-3-642-29420-4_5.

[19] Barbara Espinoza & Geoffrey Smith (2013): *Min-entropy as a resource*. *Inf. Comput.*, doi:10.1016/j.ic.2013.03.005.

[20] Yusuke Kawamoto, Konstantinos Chatzikokolakis & Catuscia Palamidessi (2014): *Compositionality Results for Quantitative Information Flow*. In: *Proc. of QEST'14*, pp. 368–383, doi:10.1007/978-3-319-10696-0_28.

[21] Yusuke Kawamoto & Thomas Given-Wilson (2015): *Quantitative Information Flow for Scheduler-Dependent Systems*. Research Report, INRIA. Available at http://hal.inria.fr/hal-01114778.

[22] Paul C Kocher (1996): *Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems*. In: *Proc. of CRYPTO'96*, Springer, pp. 104–113, doi:10.1007/3-540-68697-5_9.

[23] Boris Köpf & David A. Basin (2007): *An information-theoretic model for adaptive side-channel attacks*. In: *Proc. of CCS*, ACM, pp. 286–296, doi:10.1145/1315245.1315282.

[24] Annabelle McIver, Carroll Morgan, Geoffrey Smith, Barbara Espinoza & Larissa Meinicke (2014): *Abstract Channels and Their Robust Information-Leakage Ordering*. In: *Proc. of POST'14*, pp. 83–102, doi:10.1007/978-3-642-54792-8_5.

[25] Geoffrey Smith (2009): *On the Foundations of Quantitative Information Flow*. In: *Proc. of FOSSACS*, *LNCS* 5504, Springer, pp. 288–302, doi:10.1007/978-3-642-00596-1_21.

[26] Geoffrey Smith (2011): *Quantifying Information Flow Using Min-Entropy*. In: *Proc. of QEST'11*, pp. 159–167, doi:10.1109/QEST.2011.31.

[27] Lijun Zhang & Martin R. Neuhäußer (2010): *Model Checking Interactive Markov Chains*. In: *Proc. of TACAS'10*, Springer-Verlag, Berlin, Heidelberg, pp. 53–68, doi:10.1007/978-3-642-12002-2_5.